

# Learning Feature Relevance Through Step Size Adaptation in Temporal-Difference Learning

Alex Kearney<sup>a,b</sup>, Vivek Veeriah<sup>a,b</sup>, Jaden Travnik<sup>b</sup>, Patrick M. Pilarski<sup>b,1</sup>,  
Richard S. Sutton<sup>b,1</sup>

<sup>a</sup>*Borealis AI, Edmonton, Alberta, Canada*

<sup>b</sup>*University of Alberta, Edmonton, Alberta, Canada*

---

## Abstract

There is a long history of using meta learning as representation learning, specifically for determining the relevance of inputs. In this paper, we examine an instance of meta-learning in which feature relevance is learned by adapting step size parameters of stochastic gradient descent—building on a variety of prior work in stochastic approximation, machine learning, and artificial neural networks. In particular, we focus on stochastic meta-descent introduced in the Incremental Delta-Bar-Delta (IDBD) algorithm for setting individual step sizes for each feature of a linear function approximator. Using IDBD, a feature with large or small step sizes will have a large or small impact on generalization from training examples. As a main contribution of this work, we extend IDBD to temporal-difference (TD) learning—a form of learning which is effective in sequential, non i.i.d. problems. We derive a variety of IDBD generalizations for TD learning, demonstrating that they are able to distinguish which features are relevant and which are not. We demonstrate that TD IDBD is effective at learning feature relevance in both an idealized gridworld and a real-world robotic prediction task.

*Keywords:* Reinforcement Learning, Representation Learning, Meta-learning, Step Size Adaptation

---

## 1. Representation Learning Through Feature Relevance

The performance of Machine Learning (ML) methods depends on their representation of the environment. A representation is composed by a collection of features which describe aspects of the current state or example which is observed by the system. The choice of representation can be the difference between a successful application and one which is unable to learn. For this reason, representations are often designed by engineers who construct features which best describe the task at hand.

Effectiveness of hand-constructed of features is limited, as designing a representation requires substantial knowledge of both the environment and the problem being solved. The features which are appropriate for a given task are not necessarily transferable to different environments and problems. For each new problem and environment the engineer must re-design the representation. For this reason, hand-constructed features and engineer expertise is not a scalable approach to representation construction. Due to these limitations, it is desirable for systems to independently and automatically construct the features a system uses to learn.

The process of automatically deciding how to represent the environment can itself be learned through a second-order learning process—sometimes described as representation learning, or more broadly as both *learning to learn* and *meta-learning*. Meta-learning methods learn to modify the inputs or the parameters of the underlying machine learning method. We can use meta-learning to perform representation learning by learning to constructing new features, or learning to shaping an existing representations by weighting given features. In this paper, we focus on systems which perform meta-learning in order to shape a given representation, i.e. by identifying relevant features.

The simplest case of representation learning is learning the relevance of given features. Whether features are constructed by hand or learned, the features a system uses will not be equally relevant to the task at hand. There are some features which will be more relevant and which we wish to generalize over more than others. By learning feature relevance, a system can weight the influence of input features on the model being learnt and generalization from training examples. In this sense, learning the relevance of features is a form of representation learning which is prior to feature construction: before

creating systems which are able to construct their own features from inputs, a system can modify its existing representation by weighting given features based on their relevance. The problem of identifying relevant features through meta-learning has roots in both animal and human learning. Humans and animals are known to learn to discern which aspects of the environment are relevant to the task at hand. Work in neuroscience has assessed how humans perform representation learning by identifying relevant stimuli (Wilson and Niv, 2012); in cognitive science studies have examined how children are able to generalize from just a few examples by forming appropriate biases (Colunga and Smith, 2005). Animals learn over which features to generalize their learning to new examples—they learn the saliency of the signals. In doing so, humans and animals are performing representation learning by identifying the relevance of stimuli. It is natural that a machine learning system perform representing learning by learning the relevance of inputs.

One method of assigning feature relevance is through adapting a vector of many step sizes. Step sizes scale updates made to a learned model; by assigning step sizes on a per-feature basis we are able to give large step sizes to relevant features and small step sizes to the irrelevant features: scaling weight updates based on the quality of input features (Sutton, 1992). The simplest case of step sizes as feature relevance are methods which use a linear function approximation: methods where each weight has its own specific step size. In this paper, we focus on methods which use linear function-approximation. In the linear mapping case there is a single step size per feature and feature-relevance directly corresponds with step sizes. One of the main methods of learning feature relevance by learning setting step size values is through Stochastic Meta-descent (SMD): a form of gradient descent which takes the gradient of a gradient step.

## 2. Stochastic Meta-descent For Meta Learning And Feature Relevance

SMD was first introduced for online learning in the linear case as Incremental Delta-Bar-Delta (IDBD) (Sutton, 1992) which was later extended to non-linear mappings (Schraudolph, 1999). Most recently, SMD has been used for temporal predictions and sequential problems, such as MAML (Finn et al., 2017): an offline method for updating the step sizes in a model-agnostic

fashion. Meta-gradient RL uses SMD to adapt the return of a reinforcement learning problem (Xu et al., 2018), and Cross-prop uses SMD to learn weightings of inputs to learn representations which generalize across tasks. Outside of machine learning, IDBD has been extended to biologically plausible version for modeling neural metaplasticity (Schweighofer and Arbib, 1998).

Many generalizations of IDBD have been used to learn step sizes online for a variety of underlying learning methods. AutoStep (Mahmood and Sutton, 2013) is an extension of IDBD which introduces a normalization to the update of IDBD’s meta-weights to improve stability. SID, NOSID, Dabney AutoStep (Dabney, 2014), and our initial presentation of TIDBD (Kearney et al., 2017) are generalizations of IDBD for TD methods. We are particularly interested in these meta-learning methods for step size adaptation to learn the relevance of features.

## IDBD

Incremental Delta-Bar-Delta(IDBD) (Sutton, 1992) is a meta-learning algorithm which learns a bias through experience by maintaining a vector of learned step sizes. The intuition behind IDBD is that features which are correlated with our prediction problem should be given large step sizes, while features which are irrelevant to our prediction problem should be given smaller step sizes.

IDBD estimates a target value  $\hat{y}$  by learning a weight vector  $w$ , such that the dot-product of the observations and the weight vector produce an estimate  $y = w^\top x$ . The weights are updated by moving in the direction of the error  $\delta = \hat{y} - y$ , where  $\hat{y}$  is the observed target. The update is weighted by a small fraction  $\alpha > 0$ , where  $\alpha$  is a vector of step sizes which are learnt and  $\alpha_i > 0$ .

IDBD adapts a vector of many step sizes online and incrementally by performing stochastic meta-descent over a vector of meta-weights  $\beta$  which are used to specify the step sizes  $\alpha$  on a per-feature basis. On each time-step alpha is updated by  $\alpha \leftarrow \exp(\beta)$ , producing a vector of step sizes: one step size for each feature in observations  $x$ . By exponentiating the meta-weights  $\beta$  to produce a step size  $\alpha$ , a linear step in the meta-weights  $\beta$  produces a geometric step in  $\alpha$ ; in addition, it ensures all step sizes  $\alpha$  are positive.

The meta-weights are updated by  $\beta \leftarrow \beta + \theta\delta xh$ . The meta step size  $\theta$  weights the amount by which we update our weights in the direction of the error. The prediction error  $\delta$  is the difference between our estimate  $y$  and the observed target  $\hat{y}$ . The additional memory vector  $h$  is a decaying trace of recent meta-weight updates. The update of the  $\beta$  is proportional to the correlation between the current change and previous changes to  $\beta$ .

The intuition behind this update is that if many updates for feature  $x_i$  to weight  $w_i$  are similar, it would have been a more effective use of data to have a larger step size  $\alpha_i$  and thus a larger update to  $w_i$ . Negative correlation suggests that updates to  $w_i$  have over-shot and over-corrected—that the step size  $\alpha_i$  should be smaller.

---

**Algorithm 1** IDBD

---

- 1: Initialize vectors  $h$ ,  $\beta$ , and  $w$  of size  $n$  number of features.
  - 2: Repeat for each observation  $x$  and target  $\hat{y}$ :
  - 3:    $y \leftarrow w^\top x$
  - 4:    $\delta \leftarrow \hat{y} - y$
  - 5:   For  $i = 1, 2, \dots, n$ :
  - 6:      $\beta_i \leftarrow \beta_i + \theta\delta x_i h_i$
  - 7:      $\alpha_i \leftarrow e^{\beta_i}$
  - 8:      $w_i \leftarrow w_i + \alpha_i\delta x_i$
  - 9:      $h_i \leftarrow h_i[1 - \alpha_i x_i^2]^+ + \alpha_i\delta x_i$
- 

While it is presented alongside the underlying gradient-descent learning mechanism, IDBD is truly the update and maintenance of  $\beta$ ,  $\alpha$ , and  $h$  (lines 6,7, and 9, respectively). In this sense, it is a meta-learning method which is distinct and separate from the method which learns resulting model; however, the order of IDBD’s updates in relation to the underlying learning updates is important.

A possible criticism of IDBD is that it is only abstracting the problem of setting the step sizes to a higher level: although IDBD learns the step size parameter, we must now specify the meta step size  $\theta$  with which we learn our step sizes. This is still an improvement, as tuned IDBD outperforms methods which do not adapt their bias (Sutton, 1992). In addition, extensions of IDBD, including AutoStep (Mahmood et al., 2012), and NOSID, and AUTOSID (Dabney, 2014) have had success in making their methods relatively invariant

to the setting of the meta step size  $\theta$ .

Another criticism of IDBD is that of stability. If many consecutive weight updates to  $w_i$  are in the same direction, then  $h_i$  will grow correspondingly large. As  $h_i$  grows,  $\beta_i$  grows, and  $\alpha_i$  increases geometrically in size. This could lead to instability as a single large update to  $\beta_i$  could lead to divergence. To prevent this, Sutton (1992) originally suggests that updates to meta-weights  $\beta$  are limited  $\pm 2$  and that each step size  $\alpha_i$  is limited to some maximum value. While this resolves the instability of IDBD, it introduces two thresholds to the algorithm. This adds greater complexity to the practical use of IDBD, as we must pick appropriate thresholds for each problem. While IDBD abstracts the problem of tuning, it does not fully escape it.

### 3. Temporal Difference Learning

A limitation of IDBD is that it is derived for supervised learning, preventing its use on other learning problems which update their models with different underlying learning methods. Most existing extensions of IDBD are generalizations to other learning methods. In particular, reinforcement learning policy evaluation methods have received many generalizations: for instance, SID, NOSID, and an AutoStep variation (Dabney, 2014).

A form of policy evaluation is Temporal-difference learning (TD). TD methods are of note, as they are able to learn directly from experience, and their estimates are learned using *bootstrapping*—TD methods are able to update their estimate  $V(s)$  based on their current estimate. Temporal-difference learning methods are driven by the difference between estimates of successive states.

We can think of TD predictions as estimating the value of a state in a Markov Reward Process (MRP). A MRP is described by the tuple  $\langle S, p, r, \gamma \rangle$  where  $S$  is the set of all states,  $p(s'|s)$  describes the probability of a transition from a state  $s \in S$  to a new state  $s' \in S$ ,  $r(s, s')$  describes the reward observed on a transition from  $s$  to  $s'$ , and  $0 \leq \gamma \leq 1$  is a discount factor which determines how future reward is weighted.

The goal in an MRP is to learn a value function  $V(s)$  which estimates the expected return from a given state  $v^*(s) := \mathbb{E}\{G_t | S_t = s\}$ , where the return is  $G_t := \sum_{i=1}^{\infty} \gamma^{i-1} R_{t+i}$ , or the discounted sum of all future rewards. Within the context of an MRP, a prediction is an estimation of the value of a state—an accumulation of discounted future reward. For example, the prediction signal could be the position of a robot’s gripper, which is later used as an input into a robot’s control system.

We describe TD learning with eligibility traces and linear function approximations. The estimate  $V(s)$  is learnt by updating a vector of weights  $w$ . Using some linear function approximation—i.e.: Tile Coding (Sutton and Barto 1998) or Selective Kanerva Coding (Travník and Pilarski 2017)—we can generate a binary feature vector which represents state  $s$ . The estimated value of state  $s$  is then  $V(s) = w^\top \phi(s)$ .

The error we minimize is the TD error  $\delta = R + \gamma w^\top \phi(s') - w^\top \phi(s)$ . Note,  $R + \gamma w^\top \phi(s')$  is bootstrapped estimate of the true return  $G$ . We estimate the true return by taking the sum observed reward  $R$  and the discounted estimate of the following observed state  $V(s') = \gamma w^\top \phi(s')$ . This means that we are learning our estimate of the target  $G$  through experience. The discounting factor is  $0 \leq \gamma \leq 1$ .

The vector  $z$  is a decaying trace of recently activated features. At each time step, the trace is decayed by  $\gamma\lambda$  and then incremented by the most recently activated features. The parameter  $0 \leq \lambda \leq 1$  describes the rate at which we want to decay our traces. This enables current rewards to be attributed to previously visited states. The weights  $w$  are then incremented in the direction of the TD error  $\delta$ , weighted by a small, positive step size  $\alpha$ , and our eligibility traces  $\lambda$ .

TD prediction methods are effective in life-long continual learning systems—systems where the true return  $G_t$  may never be observed, or where the dynamics of they system are too complex to model. While useful, TD learning has received relatively little interest in step size adaptation methods. Those which almost exclusively adapt, scalar step sizes: they do not adapt bias by performing representation learning. In this paper, we provide a step size based bias adaptation method by generalizing IDBD to TD learning.

---

**Algorithm 2** TD( $\lambda$ )

---

- 1: Initialize vectors  $z \in 0^n$ , and both  $w \in \mathbb{R}^n$ ; initialize a small scalar  $\alpha > 0$ ; observe state  $S$
  - 2: Repeat for each observation  $s'$  and reward  $R$ :
  - 3:    $\delta \leftarrow R + \gamma w^\top \phi(s') - w^\top \phi(s)$
  - 4:   For  $i = 1, 2, \dots, n$ :
  - 5:      $z_i \leftarrow z_i \gamma \lambda + \phi_i(s)$
  - 6:      $w_i \leftarrow w_i + \alpha_i \delta z_i$
  - 7:    $s \leftarrow s'$
- 

#### 4. TIDBD: TD Incremental Delta-Bar-Delta

We now generalize IDBD to TD learning. There exist four other generalizations of IDBD for policy evaluation: Scalar Incremental Delta-Bar-Delta (SID), Normalized Scalar Incremental Delta-Bar-Delta (NOSID), an AutoStep variation (Dabney, 2014), and meta-trace (Young et al., 2018). Both SID and NOSID use a single, shared, global step size; as a result, neither are learning feature relevance. Although SID and NOSID learn step sizes using IDBD, the step sizes are not assigned on a per-feature basis. Dabney’s AutoStep uses a

In Section 7.1 We compare our methods against SID, NOSID, and AlphaBound (Dabney and Barto, 2012) as a sanity-check to ensure bias change and representation learning provide some benefit over methods without representation learning. In addition, we compare against (Dabney, 2014)’s AutoStep.

IDBD was originally derived as meta gradient-descent for Least Means Square rule learning. IDBD minimizes  $\frac{\partial \delta^2(t)}{\partial \beta_i(t)}$  where  $\delta$  is LMS error and  $\beta$  are the meta-weights such that  $\exp(\beta) = \alpha$ . To generalize IDBD to TD learning we must define what squared error we aim to minimize. One option is to

minimize the squared one-step TD error with respect to our meta-weights  $\beta$ .

$$\begin{aligned}\beta_i(t+1) &= \beta_i(t) - \frac{1}{2}\theta \frac{\partial \delta^2(t)}{\partial \beta_i} \\ &= \beta_i(t) - \frac{1}{2}\theta \sum_j \frac{\partial \delta^2(t)}{\partial w_j(t)} \frac{\partial w_j(t)}{\partial \beta_i}\end{aligned}\tag{1}$$

To approximate  $\sum_j \frac{\partial \delta^2(t)}{\partial w_j(t)} \frac{\partial w_j(t)}{\partial \beta_i}$  we assume that  $\frac{\partial w_j(t)}{\partial \beta_i} \approx 0$  where  $i \neq j$ . This approximation is fair, as the effect of changing the step size for a particular weight will predominantly be on the weight itself; effects on other weights will be nominal.

$$\beta_i(t+1) \approx \beta_i(t) - \frac{1}{2}\theta \frac{\partial \delta^2(t)}{\partial w_i(t)} \frac{\partial w_i(t)}{\partial \beta_i}\tag{2}$$

The use of TD error introduces some subtleties based on bootstrapping: the estimate of the error  $\delta = R_t + V(\phi(t+1)) - V(\phi(t))$  depends on the predicted value of the future state  $V(\phi(t+1))$ , resulting in a biased gradient.  $R_{t+1} + \gamma V(\phi(t+1)) - V(\phi(t))$  is a biased estimate of the expected return from time-step  $t$ , as it relies on the estimate produced by the weight vector  $w_t$ . Because of this bootstrapping, we are not using true gradient descent (Barnard, 1993).

We have two choices: performing gradient descent using the full, biased gradient, or using a semi-gradient method. Semi-gradient methods do not use the estimate of the return at state  $\phi(t+1)$  in the error. In this section, we show the derivation for both choices. In the following section, we evaluate the performance of each method.

### *Semi-gradient Derivation*

For semi-gradient TD( $\lambda$ ) we take the gradient of the approximate value function  $V$  with respect to our weight vector  $w$ . In the linear case, the

semi-gradient for TD( $\lambda$ ) is simply  $-\phi(t)$ . Using this gradient we can find (3) which may then be substituted back into (1).

$$\begin{aligned}
-\frac{1}{2} \frac{\partial \delta^2(t)}{\partial w_i(t)} &= -\delta(t) \frac{\partial \delta(t)}{\partial w_i(t)} \\
&= -\delta(t) \frac{\partial}{\partial w_i(t)} [-V(\phi(t))] \\
&= \delta(t) \phi(t)
\end{aligned} \tag{3}$$

We then complete  $\beta$ 's definition in (4) by defining an additional memory vector  $h$ . We define  $h$  as  $\frac{\partial w_i(t+1)}{\partial \beta_i}$ .

$$\beta_i(t+1) \approx \beta_i(t) + \theta \delta(t) \phi_i(t) h_i(t) \tag{4}$$

We simplify  $\frac{\partial w_i(t)}{\partial \beta_i}$  by describing it in terms of it's update.

$$\begin{aligned}
h_i(t+1) &= \frac{\partial w_i(t+1)}{\partial \beta_i} \\
&= \frac{\partial}{\partial \beta_i} [w_i(t) + e^{\beta_i(t+1)} \delta(t) z_i(t)] \\
&= h_i(t) + e^{\beta_i(t+1)} \delta(t) z_i(t) + e^{\beta_i(t+1)} \frac{\partial \delta(t)}{\partial \beta_i} z_i(t) + e^{\beta_i(t+1)} \frac{\partial z_i(t)}{\partial \beta_i} \delta_i(t)
\end{aligned} \tag{5}$$

using the product rule to simplify (5) leaves us with the remaining  $\frac{\partial \delta(t)}{\partial \beta}$ .

$$\begin{aligned} \frac{\partial \delta(t)}{\partial \beta_i} &= -\frac{\partial}{\partial \beta_i} [V(\phi(t))] \\ &= -\frac{\partial}{\partial \beta_i} \sum_j w_j(t) \phi_j(t) \\ &\approx -\frac{\partial}{\partial \beta_i} [w_i(t) \phi_i(t)] = -h_i(t) \phi_i(t) \end{aligned} \quad (6)$$

Again, as we presume that a change in step size for a particular weight will have a nominal impact on other weights, we approximate  $\frac{\partial V_i(t)}{\partial \beta_i}$  as  $\frac{\partial}{\partial \beta} [w_i(t)^\top \phi(t)]$ . This results in  $-h(t) \phi(t)$  which we may then use to simplify (5) to the definition of  $h$  in (8).

$$\frac{\partial z_i(t+1)}{\partial \beta_i} = \frac{\partial}{\partial \beta_i} [\gamma \lambda z_i(t) + \phi_i(s_t)] = \frac{\partial z_i(t) \gamma \lambda}{\partial \beta_i} = 0 \quad (7)$$

We see that (7) results in a decaying trace of the initialized value of the eligibility traces. Since the gradient is 0, this value will always be 0.

$$\begin{aligned} h_i(t+1) &\approx h_i(t) + e^{\beta_i(t+1)} \delta(t) z_i(t) - e^{\beta_i(t+1)} \phi_i(t) z_i(t) h_i(t) \\ &= h_i(t) [1 - \alpha(t+1) \phi_i(t) z_i(t)] + \alpha_i(t+1) \delta(t) z_i(t) \end{aligned} \quad (8)$$

After positively bounding  $[1 - \alpha(t+1) \phi_i(t) z_i(t)]$ , denoted with  $+$ , we have completed semi-gradient TIDBD, as shown in algorithm 3

One may note that semi-gradient TIDBD is similar to the original IDBD formulation (Algorithm 1). The meta-weights are updated using the product of the current error  $\delta$ , a trace of recent weight updates  $h$ , and the currently active features  $\phi(s)$ . The most notable change is that TIDBD's  $h$  trace is now modulated by not just the active features  $\phi$ , but also the eligibility traces  $z$ . This means that while the updates to step sizes will be limited to currently active features  $\phi(s)$ , the trace of recent weight updates include discounted past activations in  $z$ .

---

**Algorithm 3** TIDBD( $\lambda$ ) with semi-gradient

---

- 1: Initialize vectors  $h \in \mathbb{R}^n$ ,  $z \in \mathbb{R}^n$ , and both  $w \in \mathbb{R}^n$  and  $\beta \in \mathbb{R}^n$  as desired; initialize a scalar  $\theta$ ; observe state  $S$
  - 2: Repeat for each observation  $s'$  and reward  $R$ :
  - 3:    $\delta \leftarrow R + \gamma w^\top \phi(s') - w^\top \phi(s)$
  - 4:   For element  $i = 1, 2, \dots, n$ :
  - 5:      $\beta_i \leftarrow \beta_i + \theta \delta \phi_i(s) h_i$
  - 6:      $\alpha_i \leftarrow e^{\beta_i}$
  - 7:      $z_i \leftarrow z_i \gamma \lambda + \phi_i(s)$
  - 8:      $w_i \leftarrow w_i + \alpha_i \delta z_i$
  - 9:      $h_i \leftarrow h_i [1 - \alpha_i \phi_i(s) z_i]^+ + \alpha_i \delta z_i$
  - 10:    $s \leftarrow s'$
- 

*Ordinary Gradient Derivation*

We now derive TIDBD as stochastic meta-descent using the ordinary gradient. We start the derivation of TIDBD by describing the update rule for  $\beta$ —the meta-weights with which we define our step size.

Instead of using the semi-gradient, we consider both the estimated value of the state, and the target.

$$\begin{aligned} \beta_i(t+1) &\approx \beta_i(t) - \theta \delta(t) \frac{\partial \delta(t)}{\partial w_i(t)} \frac{\partial w_i(t)}{\partial \beta_i} \\ &= \beta_i(t) - \theta \delta(t) \frac{\partial [R + \gamma w^\top \phi(t+1) - w^\top \phi(t)]}{\partial w_i(t)} \frac{\partial w_i(t)}{\partial \beta_i} \quad (9) \\ &= \beta_i(t) - \theta \delta(t) [\gamma \phi(t+1) - \phi(t)] \frac{\partial w_i(t)}{\partial \beta_i} \\ &= \beta_i(t) - \theta \delta(t) [\gamma \phi(t+1) - \phi(t)] h_i(t) \end{aligned}$$

We then complete the simplification of  $\beta$ 's update by defining an additional

memory vector  $h$  as  $\frac{\partial w}{\partial \beta}$ . We then complete the update for  $h$ .

$$\begin{aligned}
h_i(t+1) &= \frac{\partial w_i(t+1)}{\partial \beta_i} \\
&= \frac{\partial [w_i(t) + e^{\beta_i(t+1)} \delta(t) z_i(t)]}{\partial \beta_i} \\
&= h_i(t) + e^{\beta_i(t+1)} \delta(t) z_i(t) + e^{\beta_i(t+1)} \frac{\partial \delta(t)}{\partial \beta_i} z_i(t) + e^{\beta_i(t+1)} \frac{\partial z_i(t)}{\partial \beta_i} \delta_i(t)
\end{aligned} \tag{10}$$

This simplification leaves us with  $\frac{\partial \delta(t)}{\partial \beta_i}$ , derived in (6), and  $\frac{\partial z_i(t)}{\partial \beta_i}$ , simplified in (7). We use the same approximation as in (1) to simplify:

$$\begin{aligned}
\frac{\partial \delta(t)}{\partial \beta_i} &= \frac{\partial}{\partial \beta_i} [R + \gamma w^\top \phi(t+1) - w^\top \phi(t)] \\
&= \frac{\partial}{\partial \beta_i} \left[ \sum_j R + \gamma w_j \phi_j(t+1) - w_j \phi_j(t) \right] \\
&\approx \frac{\partial}{\partial \beta_i} [R + \gamma w_i \phi_i(t+1) - w_i \phi_i(t)] \\
&= \gamma h_i \phi_i(t+1) - h_i \phi_i(t)
\end{aligned} \tag{11}$$

$$\begin{aligned}
h_i(t+1) &\approx h_i(t) + e^{\beta_i(t+1)} \delta(t) z_i(t) + e^{\beta_i(t+1)} [\gamma h_i \phi_i(s_{t+1}) - h_i \phi_i(s_t)] z_i(t) + 0 e^{\beta_i(t+1)} \delta_i(t) \\
&= h_i(t) [1 + \alpha_i(t+1) z_i(t) [\gamma \phi_i(s_{t+1}) - \phi_i(s_t)]] + \alpha_i(t+1) \delta(t) z_i(t)
\end{aligned} \tag{12}$$

We then take the results from (6) and (7) to complete the definition of  $h$ 's update. The update for  $h$  and  $\beta$  may then be implemented directly as shown earlier in Algorithm 4

---

**Algorithm 4** TIDBD( $\lambda$ )

---

- 1: Initialize vectors  $h \in \mathbb{R}^n$ ,  $z \in \mathbb{R}^n$ , and both  $w \in \mathbb{R}^n$  and  $\beta \in \mathbb{R}^n$  as desired; initialize a scalar  $\theta$ ; observe state  $S$
  - 2: Repeat for each observation  $s'$  and reward  $R$ :
  - 3:    $\delta \leftarrow R + \gamma w^\top \phi(s') - w^\top \phi(s)$
  - 4:   For element  $i = 1, 2, \dots, n$ :
  - 5:      $\beta_i \leftarrow \beta_i - \theta \delta [\gamma \phi(s') - \phi(s)] h_i$
  - 6:      $\alpha_i \leftarrow e^{\beta_i}$
  - 7:      $z_i \leftarrow z_i \gamma \lambda + \phi_i(s)$
  - 8:      $w_i \leftarrow w_i + \alpha_i \delta z_i$
  - 9:      $h_i \leftarrow h_i [1 + \alpha_i z_i [\gamma \phi_i(s') - \phi_i(s)]]^+ + \alpha_i \delta z_i$
  - 10:   $s \leftarrow s'$
- 

## 5. Does TIDBD(0) With a Single, Shared step size Outperform Ordinary TD?

Having derived both ordinary and semi-gradient TIDBD, we now evaluate whether the benefits of IDBD transfer to TD learning.

First, we assess the ability of TIDBD to improve upon traditional TD prediction in a simple tabular setting. In a tabular setting, the advantages of vectorizing step sizes are not present, enabling us to assess whether adapting step sizes with TIDBD is an improvement over ordinary TD in general, independent of performing representation learning.

An ideal bias learning method would be able to perform as well as or better than TD for arbitrary initial step sizes  $\alpha_0$  while being insensitive to meta-parameters. We expect that TIDBD should be able to out-perform TD for all  $\theta$  values, and that it should be relatively insensitive to the choice of  $\theta$  values.

### 5.1. Gridworld

We created a suitable prediction task by generating a Markov Reward Process from a grid-world problem originally described in [Sutton and Barto \(1998\)](#) (depicted in Figure [1](#)). Each tile in the  $5 \times 5$  grid-world represents a

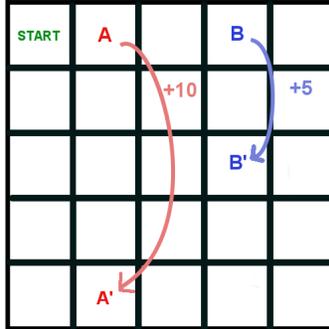


Figure 1: Gridworld Problem as introduced in (Sutton and Barto, 1998).

state. The state transitions are the four cardinal directions—north, south, east, and west—chosen by an equiprobable random policy. Transitions which would leave the grid resulted in staying in the same state and a reward of -1. Regardless of the transition in state  $A$  or  $B$ , the learner transitions to states  $A'$  and  $B'$  respectively with a probability of 1. A transition from  $A$  to  $A'$  yields a reward of 10 and a transition from  $B$  to  $B'$  yields a reward of 5. All other transitions receive a reward of 0. The start state was the top left-hand corner. A trial consisted of the equiprobable random policy acting for 15000 time-steps. Each prediction method would then learn a value function over the 30 trials.

We compared ordinary TD to both semi-gradient and ordinary gradient TIDBD with initial step sizes distributed between 0.0005 and 0.5. For all prediction methods  $\lambda = 0$  and  $\gamma = 0.99$ . We swept over 21 different meta-parameters equally distributed between the range of  $0 < \theta < 0.2$ —the range for which IDBD was originally compared over in Sutton (1992). When  $\theta = 0$ , TIDBD and TD are equivalent, as the initial step size  $\alpha_0$  is never updated.

Figure 2 depicts the performance of both semi-gradient and ordinary gradient TIDBD for settings of their meta step size. As expected, semi-gradient TIDBD is less sensitive to meta step size  $\theta$  values than ordinary TIDBD, but has a higher asymptotic error than ordinary TD for the best initial step size setting  $\alpha_0 = 0.05$ . For all but  $\alpha_0 = 0.05$ , there are broad settings of  $\theta$  such that TIDBD attains better asymptotic performance than ordinary TD. As we move further away from the optimal step size, the greater

the advantage of adapting the step sizes with TIDBD becomes.

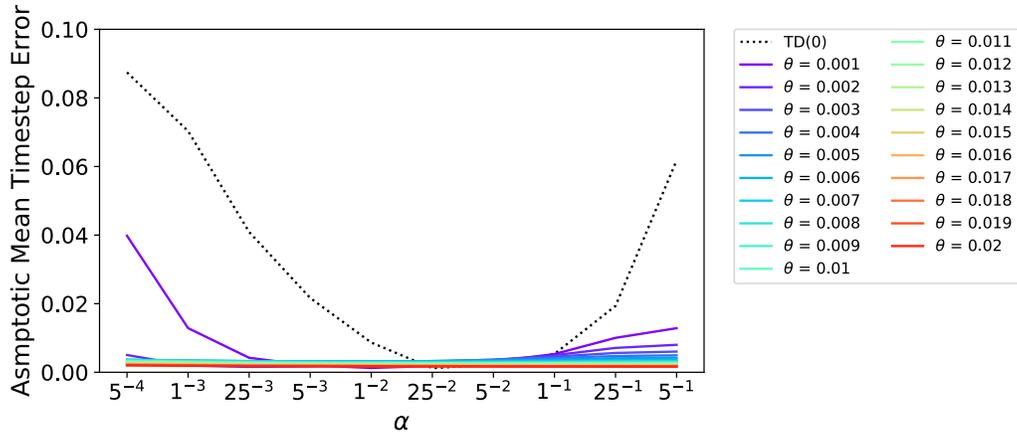
The semi-gradient accounts for the effect of changing the weights on the estimate, but ignores the effects on the target. While semi-gradient methods do not converge as robustly, they converge fast with reliability. The motivation for a semi-gradient TIDBD is to make the method less sensitive to the setting of  $\theta$ . By not taking into account the target, non-stationarity in the update of the target value will not affect the updates of, making TIDBD more stable than an ordinary gradient TD method.

For every initial step size except for the best setting ( $\alpha_0 = 0.05$ ), there are settings of  $\theta$  such that TIDBD is an improvement over ordinary fixed step size TD. Common to both ordinary and semi-gradient TIDBD is the asymmetry of sensitivity to meta-parameters: the performance of TIDBD is more robust for initializations of  $\alpha > 0.05$ . In general, TIDBD is more robust to larger initial step size settings than it is to smaller initial step size settings. This is intuitive. Step size schedules typically start with large values which decrease over time.

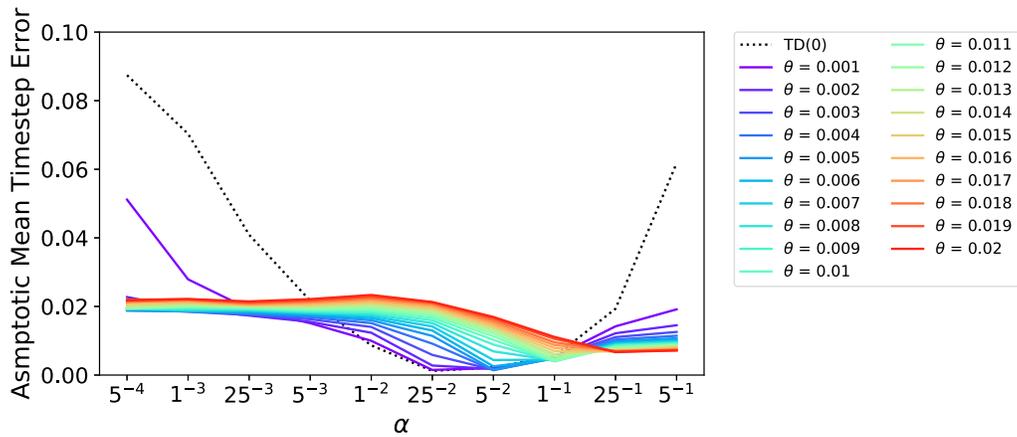
In figure [2b](#), the performance of TIDBD with an ordinary gradient is displayed across varying settings of  $\alpha$  and  $\theta$ . As anticipated, TIDBD with an ordinary gradient is more sensitive to  $\theta$  settings. Unlike semi-gradient TIDBD, TIDBD with an ordinary TD gradient achieves performance equivalent to ordinary TD for the best initial step size.

In Figure [3](#) the error of both semi-gradient and ordinary TIDBD are compared to TD for three initial step size settings. Figure [3a](#) is the most conservative initialization of  $\alpha_0$ . We expect that for this  $\alpha_0$  learning for ordinary TD will be slow, as it takes small steps when updating its estimated value for each state. For  $\alpha_0 = 0.01$ , both ordinary and semi-gradient TIDBD outperform ordinary TD, with both increasing their step sizes enabling them to make larger updates to their weights and learn more quickly.

For the best setting  $\alpha_0 = 0.05$  (Figure [3b](#)), there is little asymptotic difference between TD both versions of TIDBD; however, semi-gradient TIDBD learns far faster than both ordinary TD and ordinary gradient TIDBD. In early learning, semi-gradient TIDBD is able to find step sizes which enable it to converge the fastest. The benefits of both semi and ordinary-gradient

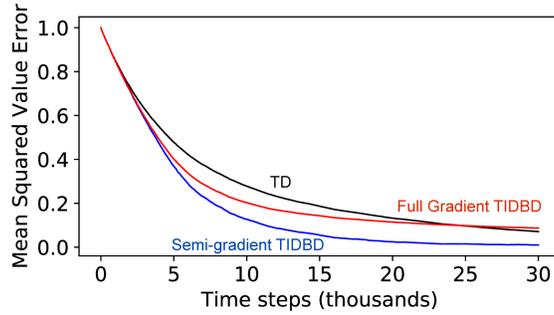


(a) Semi-gradient TIDBD.

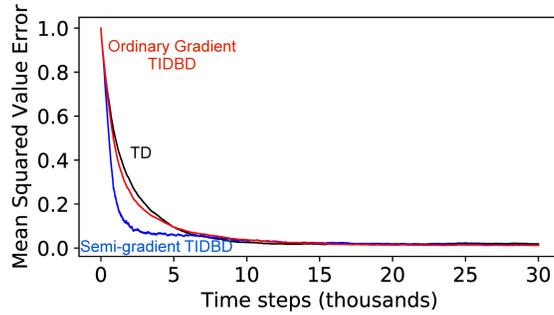


(b) Ordinary gradient TIDBD.

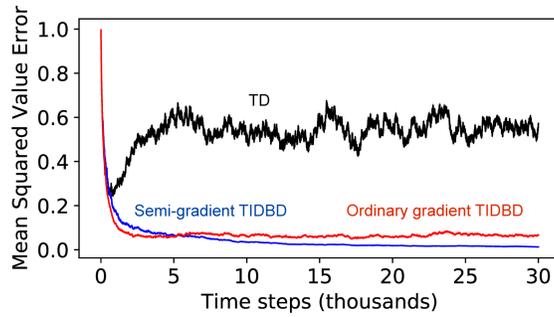
Figure 2: Parameter study of semi and ordinary gradient TIDBD. Static step size TD in black.



(a)  $\alpha_0 = 0.01$



(b)  $\alpha_0 = 0.05$



(c)  $\alpha_0 = 0.5$

Figure 3: Error during learning for three initializations  $\alpha_0$  for TD(0) and both semi-gradient TIDBD and ordinary-gradient TIDBD.

TIDBD are most apparent at the most aggressive setting of  $\alpha_0 = 0.5$  (Figure 3c). While ordinary TD is unable to learn, both semi and ordinary gradient TIDBD are able to learn in spite of the poor initial step size setting—they are able to tune their step sizes down in response.

## 6. AutoTIDBD: AutoStep-Style TIDBD

TIDBD with an ordinary gradient was able to perform as well as or better than TD for some  $\theta$  value at each  $\alpha_0$  value; however, the best  $\theta$  values varied for different  $\alpha_0$  values. Conversely, semi-gradient TIDBD had broad ranges of values of  $\theta$  for which the performance was acceptable, but did not perform as well as TD for the best value of  $\alpha_0$ .

One of the benefits of a static shared step size is that most reinforcement learning practitioners have an intuition of what range of values will yield acceptable performance in general; however, the optimal step size value will vary from domain to domain. With TIDBD and other similar IDBD methods, the performance is dependent on a meta step size—a value for which there is little intuition as to how to set. A poor choice of meta step size can lead to explosive updates.

The instability of IDBD is not a new observation. Sutton (1992) originally suggested bounding both step size values and the size of updates to the meta-weights  $\beta$ —amendments we chose to exclude so as to better understand TIDBD’s underlying performance. AutoStep—an extension of IDBD—reduced sensitivity to the setting of the meta step size and prevented divergence in the supervised learning setting (Mahmood et al., 2012). To prevent divergence, AutoStep makes two additions to IDBD. First, the meta-weight update is normalized by a decaying average of recent weight updates. Second, the step sizes is normalized by the amount by which the error was reduced on a given example—termed the *effective step size*. By normalizing the current step sizes by the effective step size a weight update will never overshoot on the current observed example.

### 6.1. *AutoStep*

To manage explosive growth of step sizes, Autostep (Algorithm 5) adds two components to IDBD (Algorithm 1): a normalization of the inputs  $\delta xh$  with which we update our  $\beta$  meta-weights, and a normalization of the resulting step size  $\exp(\beta)$  by the *effective step size*.

To normalize the inputs, AutoStep maintains  $\eta$ , a running trace of recent weight updates (Line 5). At each time-step it takes the maximum between the current weight update  $|\delta xh|$  and a decay of the previous maximum  $v_i + \frac{1}{\tau}\alpha_i x_i^2 (|\delta x_i h_i| - \eta)$ . The scalar  $\tau$  is a large value which weights the decay of  $v$ .

$$\eta \leftarrow \max(|\delta xh|, \eta + \frac{1}{\tau}\alpha x^2 (|\delta xh| - \eta)) \quad (13)$$

One might consider why the maximum is decayed rather than simply stored—as is done with NOSID (see Dabney (2014) for further details). In real-world data sources, noise and other outliers could distort the absolute recorded maximum, making the normalizer adjust input values into an unrepresentative range. By decaying the maximum, we enable the system to recover gracefully and gradually from such extreme data.

On line 7, this normalizer  $\eta$  is used to make the update  $\delta xh$  unitless.

After the meta-weights have been updated, the resulting step size  $\alpha$  is normalized by the *effective step size* (line 8). The effective step size describes the amount by which we reduce the error on the current example by making a weight update. An effective step size equal to one means that the error has been entirely reduced for the current example. By dividing the current step size  $\alpha$  by  $\max(1, \text{effective step size})$ , we prevent over-shooting on a given example: we prevent an update which introduces more error.

### 6.2. *AutoTIDBD: AutoStep for TD learning*

Having introduced AutoStep, we now add Auto-step’s normalization to TIDBD to improve its stability. AutoStep prevents divergence by normalizing

---

**Algorithm 5** AutoStep

---

- 1: Initialize vectors  $h \in 0^n$ ,  $z \in 0^n$ , and both  $w \in \mathbb{R}^n$  and  $\beta \in \mathbb{R}^n$  as desired; initialize vector  $\alpha \leftarrow \exp(\beta)$  initialize a scalar  $\theta$ ; observe state  $S$
  - 2: Repeat for each observation  $x$  and target  $y$ :
  - 3:    $\delta \leftarrow y - \mathbf{w}^\top x$
  - 4:   For element  $i = 1, 2, \dots, n$ :
  - 5:      $\eta_i \leftarrow \max(|\delta x_i h_i|, \eta_i + \frac{1}{\tau} \alpha_i x_i^2 (|\delta x_i h_i| - \eta_i))$
  - 6:     If  $\eta_i \neq 0$ :
  - 7:        $\alpha_i \leftarrow \alpha_i \exp(\mu \frac{\delta x_i h_i}{\eta_i})$
  - 8:      $M \leftarrow \max(\sum(\alpha_i x_i^2), 1)$
  - 9:     For element  $i = 1, 2, \dots, n$ :
  - 10:        $\alpha_i \leftarrow \frac{\alpha_i}{M}$
  - 11:        $w_i \leftarrow w_i + \alpha_i \delta x_i$
  - 12:        $h_i(1 - \alpha_i x_i^2) + \alpha_i \delta x_i$
- 

the current step size  $\alpha_i$  by the effective step size  $\alpha_i^\top x_i^2$ —the amount by which the error on the current example is reduced by updating the weight vector. If the effective step size is one, then we have reduced all error on the given example; If the effective step size is greater than one, then we have over-corrected on a given example. If we divide the current step size by the effective step size before performing a weight update in this instance, we ensure that we do not overshoot on the given example.

We calculate the effective step size by taking the difference between the error before the weight update  $\delta_t(t)$  and the error after the weights have been updated  $\delta_{t+1}(t)$ , or  $\frac{\delta_t(t) - \delta_{t+1}(t)}{\delta_t(t)}$ . We calculate the error  $\delta_{t+1}(t)$  using the weights from time-step  $t + 1$  and the observation from time-step  $t$ .

For supervised learning, the notion of an effective step size is straightforward: there is a known target value, so the error reduced on a given time-step is directly observable. However, as previously mentioned, TD learning uses bootstrapping. For TD learning the effective step size is not an exact value, but a biased estimation dependent on how accurate the value-function is in

estimating the value of the following step.

$$\begin{aligned}
\frac{\delta_t(t) - \delta_{t+1}(t)}{\delta_t(t)} &= \frac{[R_{t+1} + \gamma V_t(\phi(t+1)) - V_t(\phi(t))] - [R_{t+1} + \gamma V_{t+1}(\phi(t+1)) - V_{t+1}(\phi(t))]}{\delta(t)} \\
&= \frac{[\gamma V_t(\phi(t+1)) - V_t(\phi(t))]}{\delta(t)} \\
&= \frac{[\gamma(V_t(\phi(t+1)) + (\alpha\delta z)^\top \phi(t+1)) - (V_t(\phi(t)) + (\alpha\delta z)^\top \phi(t))]}{\delta(t)}
\end{aligned} \tag{14}$$

We expand  $\delta_{t+1}(t)$ , as the TD error of the current time-step  $t$  using the value-functions from the following time-step,  $V_{t+1}$ . Value functions may be written recursively as the sum of the previous time-step's value-function  $V_t(\phi(s_t))$  and the current weight update  $\alpha_t \delta_t z_t$ . So,  $V_{t+1}(\phi(s_t)) = V_t(\phi(s_t)) + [\alpha_t \delta_t z_t]^\top \phi(s_t)$ .

$$\begin{aligned}
\frac{\delta_t(t) - \delta_{t+1}(t)}{\delta_t(t)} &= \frac{[\gamma V_t(\phi(t+1)) - \gamma V_t(\phi(t+1))]}{\delta(t)} \\
&= \frac{[V_t(\phi(t)) - V_t(\phi(t))] - [\gamma\alpha\delta z^\top \phi(t+1) - \alpha\delta z^\top \phi(t)]}{\delta(t)} \\
&= \frac{-[\gamma\alpha\delta z^\top \phi(t+1)] - \alpha\delta z^\top \phi(t)}{\delta(t)} \\
&= -(\alpha z)^\top [\gamma\phi(t+1) - \phi(t)]
\end{aligned} \tag{15}$$

The resulting effective step size is  $-(\alpha z)^\top [\gamma\phi(t+1) - \phi(t)]$ . This is an intuitive result, as the amount by which we will reduce our error on a given example is the difference between the the update made to the features active in the target  $\phi(t+1)$  and the changes made to the features in the state who's value we are currently estimating  $\phi(t)$ .

With the effective step size defined, what remains in defining an AutoStep for TIDBD is the weight-update's normalizing term. AutoStep simply maintains a running trace of the absolute value of the weight-updates

$\max(|\delta x_i h_i|, v_i + \frac{1}{\tau} \alpha_i x_i^2 (|\delta x_i h_i| - v))$ . The absolute weight update for TIDBD is  $|\delta[\gamma\phi(s_{t+1}) - \phi(s)]h|$ , and the current active step size is  $\alpha[\gamma\phi(s_{t+1}) - \phi(s_t)]$ . Thus, the trace  $\eta$  of the maximum weight update would be  $\max(|\delta[\gamma\phi_i(s_{t+1}) - \phi_i(s)]h_i|, \eta_i - \frac{1}{\tau} \alpha_i [\gamma\phi_i(s_{t+1}) - \phi_i(s)] z_i (|\delta\phi_i(s)h_i| - \eta_i))$

---

**Algorithm 6** AutoStep Style Normalized TIDBD( $\lambda$ )

---

- 1: Initialize vectors  $h \in 0^n$ ,  $z \in 0^n$ , and both  $w \in \mathbb{R}^n$  and  $\beta \in \mathbb{R}^n$  as desired; initialize a scalar  $\theta$ ; observe state  $S$
  - 2: Repeat for each observation  $s'$  and reward  $R$ :
  - 3:      $\delta \leftarrow R + \gamma w^\top \phi(s') - w^\top \phi(s)$
  - 4:     For element  $i = 1, 2, \dots, n$ :
  - 5:          $\eta_i \leftarrow \max[$
  - 6:              $|\delta[\gamma\phi_i(s') - \phi_i(s)]h_i|,$
  - 7:              $\eta_i - \frac{1}{\tau} \alpha_i [\gamma\phi_i(s') - \phi_i(s)] z_i (|\delta\phi_i(s)h_i| - \eta_i)]$
  - 8:     For element  $i = 1, 2, \dots, n$ :
  - 9:          $\beta_i \leftarrow \beta_i - \theta \frac{1}{\eta_i} \delta [\gamma\phi_i(s') - \phi_i(s)] h_i$
  - 10:          $M \leftarrow \max(-e^{\beta_i} [\gamma\phi_i(s') - \phi_i(s)]^\top z_i, 1)$
  - 11:          $\beta_i \leftarrow \beta_i - \log(M)$
  - 12:          $\alpha_i \leftarrow e^{\beta_i}$
  - 13:          $z_i \leftarrow z_i \gamma \lambda + \phi_i(s)$
  - 14:          $w_i \leftarrow w_i + \alpha_i \delta z_i$
  - 15:          $h_i \leftarrow h_i [1 + \alpha_i [\gamma\phi_i(s') - \phi_i(s)] z_i]^+ + \alpha_i \delta z_i$
  - 16:      $s \leftarrow s'$
- 

With the generalization to AutoTIDBD, we assess it's performance to determine whether it is able to perform better or equal to than tuned ordinary TD(0) while being relatively insensitive to its meta step size  $\theta$ , meeting one of our core criteria for an adaptive bias algorithm.

### 6.3. AutoTIDBD in Gridworld

In Figure 5 and 4, a parameter study of AutoTIDBD’s sensitivity on the task introduced in section 5.1 is presented. As with the previous experients, AutoTIDBD is adapting a single, shared step size.

Similar to semi-gradient TIDBD, there are broad ranges of  $\theta$  values for which AutoTIDBD outperforms ordinary TD; like ordinary-gradient TIDBD, AutoTIDBD performs as well as or better than TD, even for the best  $\alpha_0$  setting. While the absolute best performance may vary for different values of  $\alpha_0$ , the change in performance as  $\theta$  varies is predictable and consistent.

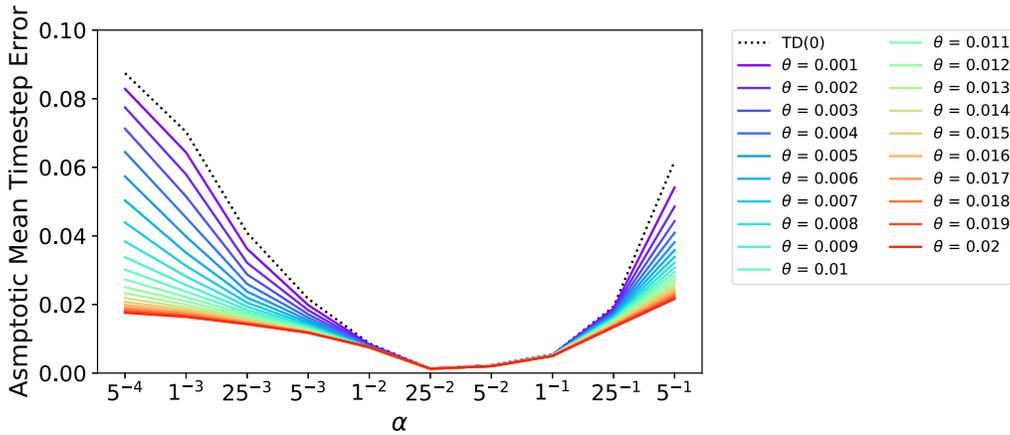


Figure 4: Parameter study of Full-gradient AutoTIDBD for varying  $\alpha_0$  and  $\theta$  values.

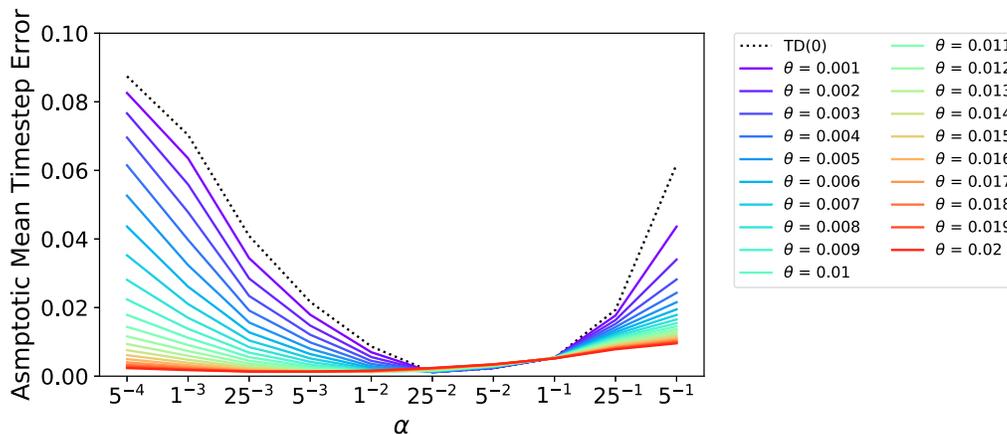
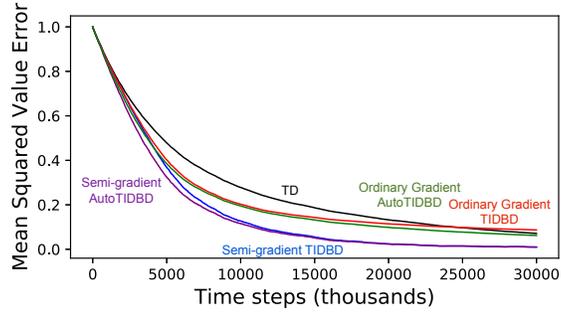
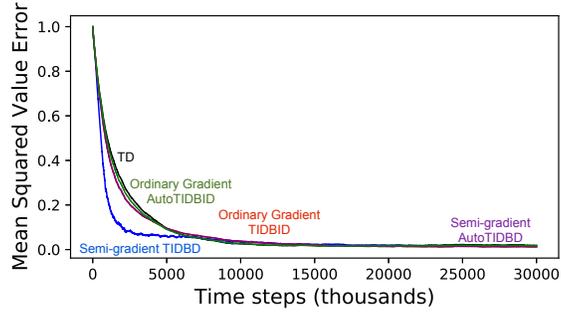


Figure 5: Parameter study of semi-gradient AutoTIDBD for varying  $\alpha_0$  and  $\theta$  values.

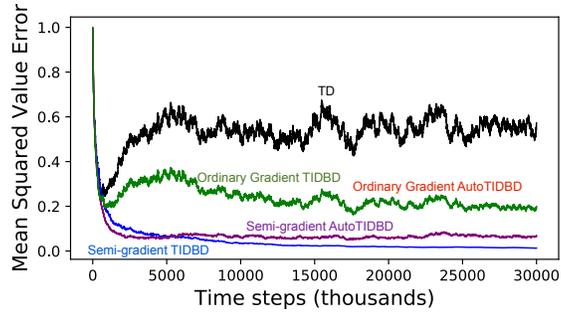
In Figure 6, the average error over the trial is shown for  $\alpha_0$  of 0.01, 0.05, and 0.5: points corresponding to initializations of  $\alpha_0$  as shown in Figures 4 and 5. We can see that for the more conservative initialization of  $\alpha_0 = 0.01$ , AutoTIDBD is able to increase its step size and learn faster than ordinary TD. As expected, AutoTIDBD performs about as well as ordinary gradient TIDBD without Auto normalization. When we start with a conservative step size, we are unlikely to diverge; AutoTIDBD and TIDBD will behave similarly in this setting. For the best initialization  $\alpha_0 = 0.05$ , AutoTIDBD performs as well as fixed step size TD. For the most aggressive setting  $\alpha_0 = 0.5$ , AutoTIDBD learns slower and more erratically than un-normalized TD.



(a)  $\alpha_0 = 0.01$

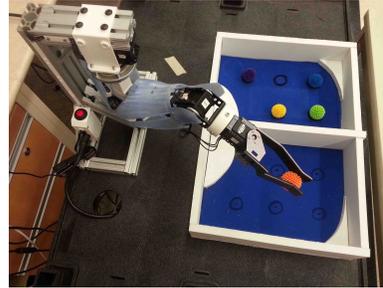
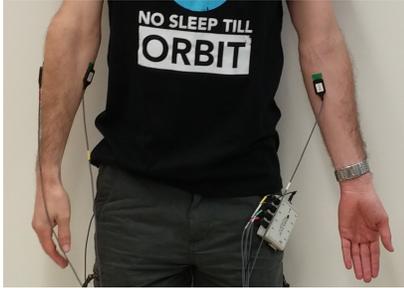


(b)  $\alpha_0 = 0.05$



(c)  $\alpha_0 = 0.5$

Figure 6: Error during learning for three initializations of  $\alpha_0$  for TD(0), both semi-gradient and ordinary-gradient TIDBD(0), and AutoTIDBD(0).



(a) A subject with electrodes attached to their wrist flexors and extensors. (b) The BentoArm performing a modified Box and Blocks task.

Unlike the previous versions of TIDBD, AutoTIDBD performs as well as or better than ordinary fixed step size TD for all initial settings of  $\alpha$ . While this stability and improved performance comes at a cost—AutoTIDBD learns more slowly for aggressive  $\alpha_0$  than TIDBD—even without representation learning AutoTIDBD is able to perform as well as or better than ordinary TD learning for all of the initial step sizes  $\alpha_0$  in the grid-world experiment.

## 7. How Robust is AutoTIDBD to Selection of Meta Step Size $\theta$

In previous sections, we demonstrated that TIDBD and AutoTIDBD tuning a single, shared step size was able to perform as well as or better than tuned ordinary TD. We now evaluate how well AutoTIDBD performs when using a vector of many step sizes—when it is performing representation learning. We evaluate this on a known, challenging, real-world prediction of problem (Pilarski et al., 2012, 2013; Seijen and Sutton, 2014).

### 7.1. Robotic Prediction Task

The prediction problem for our evaluation consisted of predicting the temporally extended future values of signals of interest within the data stream of a robotic arm as a user controlled it to perform a manual manipulation task. The dataset for this evaluation was drawn from a prior study by Edwards et al. (2016), as done in the evaluation of True Online TD by Seijen and Sutton (2014). In this dataset, the signal space of a robotic arm—the BentoArm of Dawson et al. (2014), depicted in Figure 7b—was recorded as four participants

used signals from their upper-arm muscles to control the robot to perform a simple object placement task. Signals in the data stream included the moment-by-moment position, velocity, load, temperature of all the robot’s motors, along with the control signals being sent by the human user. Users were tasked with switching between the multiple controllable degrees of the robot arm to move balls from one side of a divided box to another. For the full experimental protocol used in generating this dataset, please refer to [Edwards et al. \(2016\)](#). Important to our present evaluations, in this dataset the four participants each performed the manipulation task a total of six times: three times with a non-adaptive control system, and three times with an adaptive control system, creating a total of 24 independent time-series trials within the dataset. The inclusion of data from the adaptive control case of [Edwards et al. \(2016\)](#) adds non-stationarity to an already non-stationary prediction problem. Non-stationarity is introduced as the user becomes more proficient at the manipulation task, and as the control system begins to adapt to the user’s preferred movements. No single, scalar step size would be ideal in this setting at all times, as the prediction problem changes over time. As noted by [Pilarski et al. \(2013\)](#), finding appropriate features and setting appropriate parameters for learning systems is in fact a known challenge in this particular robotic control domain—end-user time is precious, and designers cannot possibly test their prediction algorithms on datasets which are representative of all the situations the robot might encounter upon deployment. Aspects such as non-stationarity and the irregularities introduced through human-in-the-loop control therefore make this dataset an appropriate one for studying the robustness of AutoTIDBD in predicting a real-world data stream.

### *7.2. Sensitivity to Meta step size $\theta$ or Prosthetic Prediction Problem*

For AutoTIDBD to be an improvement over ordinary, fixed step size TD, it should be less sensitive to settings of  $\theta$  than TD is to settings of  $\alpha$ : less tuning should be required for AutoTIDBD.

We constructed a prediction problem where each algorithm predicts a signal of interest from the robot arm. each algorithm predicted the angular position of the robot’s hand motor (the gripper’s aperture), as in [Seijen and Sutton \(2014\)](#). We used tile-coding to construct a binary feature vector of size  $2^{10}$  with 8 tilings and used the velocity of the hand, the position of the

hand, and the participant’s control signals to construct the feature vector. An additional bias feature was concatenated to the feature-vector, resulting in 9 active features at any given time.

We compared AutoTIDBD to NoSID, SID, AutoSID, RMSProp, and AlphaBound, as described in source papers. Each learning method used a discounting factor  $\gamma = 0.95$ . The IDBD-based adaptive step size methods shown all initialized their vector step sizes to an initial value of  $\log(\frac{1}{9})$ , which when exponentiated results in a step size  $\alpha_0 = \frac{1}{9}$ . This step size was chosen, as it was one of the best performing step size values for smaller values of  $\lambda$ , but diverges at  $\lambda = 0.6$  for ordinary fixed shared step size TD( $\lambda$ ) (shown in Figure 9). AlphaBound is initialized with a step size of 1, as originally specified.

In figure 9 the sensitivity of static step size TD( $\lambda$ ) is shown for a variety of  $\alpha$  values across  $\lambda$  settings. There is a trade-off between error and magnitude of step sizes. Consequently, there is no single step size  $\alpha$  which performs well for all  $\lambda$  values.

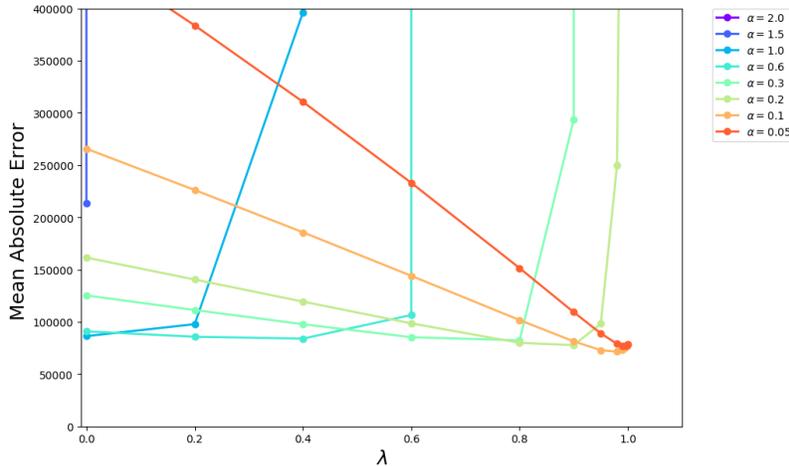


Figure 9: Average cumulative error of TD( $\lambda$ ) for various  $\alpha$  settings.

In Figure 10 the sensitivity to the same prediction problem is shown. AutoTIDBD is less sensitive to  $\theta$  than TD( $\lambda$ ) is to the setting of  $\alpha$ ; AutoTIDBD performs well for each  $\theta$  and  $\lambda$  value. This satisfies our criteria for a method which performs acceptably for broad meta-parameter settings.

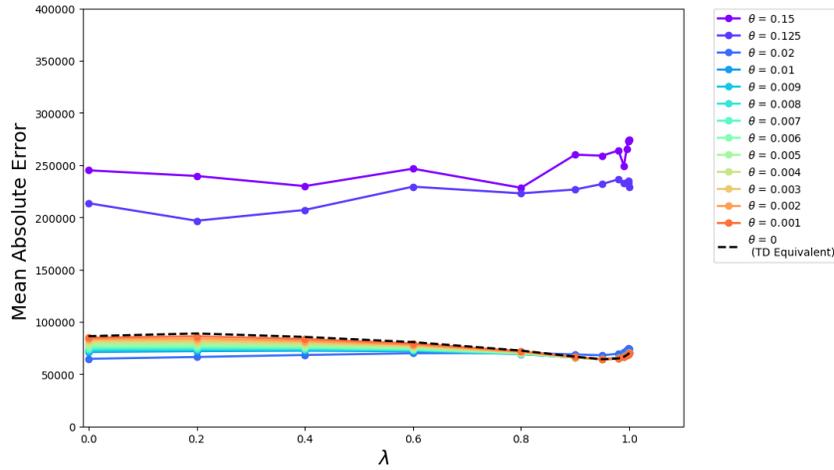


Figure 10: Average cumulative return error of AutoTIDBD for various  $\theta$  settings.

In Figure [11](#) the average cumulative error for the best tuned parameter settings for each method are shown. Adaptive step size we compared against are SID and Alphabound. These methods act as a baseline comparison. If TIDBD and AutoTIDBD are performing representation learning well, we expect them to perform better than scalar step size adaptation methods. TD with RMSProp, AutoSID, and NOSID could not be compared, as both had errors which were too large to be compared.

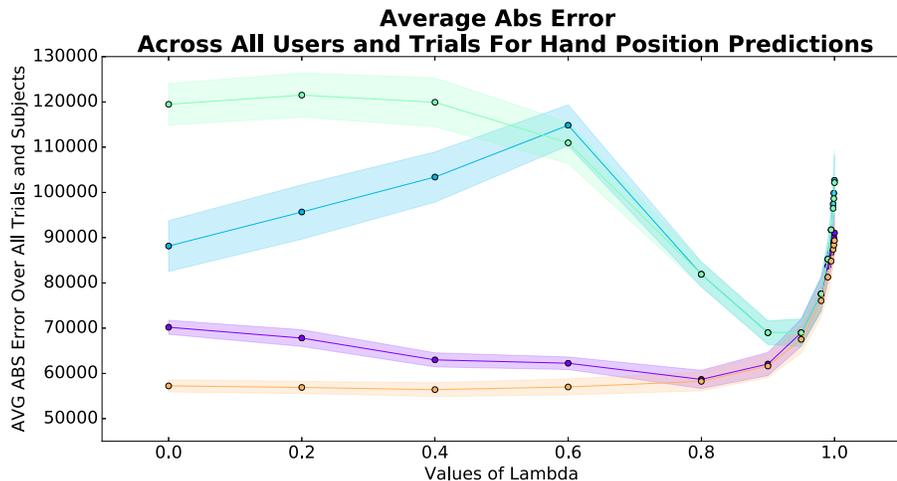


Figure 11: Absolute cumulative return error averaged over 24 independent trials. Each algorithm is presented for the best setting of  $\alpha$  or  $\theta$  and is compared by varying  $\lambda$  between 0 and 0.9

Both Semi-gradient TIDBD and Ordinary gradient TIDBD outperform SID—A scalar version of IDBD for TD learning; however, neither consistently attain errors less than or equal to ordinary TD. In contrast, AutoTIDBD attains lower or equivalent error to TD—outperforming all other adaptive methods we compared against, excluding AlphaBound for large values of  $\lambda$ .

### 7.3. Sensitivity to Meta step size $\theta$ across Prediction Problems

We previously assessed the performance of AutoTIDBD across meta step size settings for predictions of the gripper. An ideal adaptive step size method should have broad ranges of meta-parameter settings for which it attains acceptable performance. Moreover, it is ideal if these meta-parameter settings are invariant over different problems: if an ideal step size on one problem is ideal for all other problems.

Using the experiment and setup from Section 7, we compare the sensitivity of AutoTIDBD and TIDBD with an ordinary gradient for a variety of prediction formulations. We predict the velocity, position, and load for all five servos of the robot arm. Each of the signals produced by the arm are unique; as a result, each of the prediction problems is unique.

In Figure [12](#), the sensitivity of ordinary gradient TIDBD to its meta-parameter settings is depicted. Each line represents a prediction of a different signal of interest from the arm. We can see that the sensitivity for each problem is different and the best setting of  $\theta$  for each problem are different. This is problematic—the problem of tuning step sizes has simply been abstracted away to a higher-level. TIDBD may be less sensitive to initialization of its parameters than ordinary TD, but it is still sensitive and thus must be tuned for each target domain.

In Figure [13](#) the same meta-sensitivity for AutoTIDBD is depicted for the same meta-parameter settings. Each of the settings eventually diverge, but the valley of  $\theta$  parameters for which we attain reasonable performance is far broader than TIDBD and TD( $\lambda$ ), and relatively invariant across prediction problems.

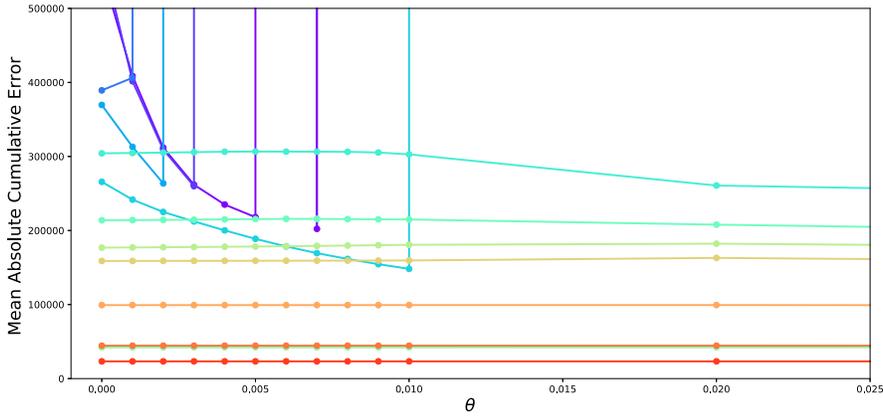


Figure 12: Average absolute cumulative return error of TIDBD(0) for different values of  $\theta$ . Each line represents the error of a prediction with its own unique signal of interest. Each on-policy prediction has the same  $\gamma$ , but is predicting a different signal of interest.

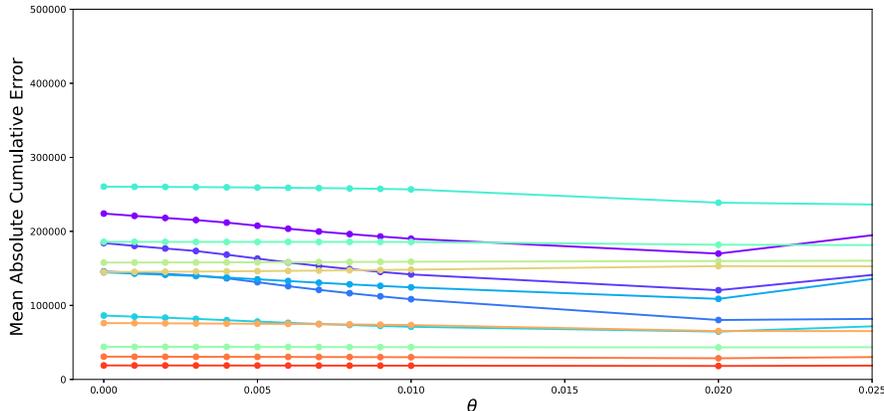


Figure 13: Average absolute cumulative return error of TIDBD(0) for different values of  $\theta$ . Each line represents the error of a prediction with its own unique signal of interest. Each on-policy prediction has the same  $\gamma$ , but is predicting a different signal of interest.

AutoTIDBD meets our requirements: It performs as well as or better than ordinary  $\text{TD}(\lambda)$ —even when only adapting a single step size, it is less sensitive to settings of  $\theta$  than  $\text{TD}(\lambda)$  is to  $\alpha$ , and has broad ranges of  $\theta$  which are invariant across problems. AutoTIDBD is bias-adaptation step size method which does not require tuning across applications. What remains of our criteria to be evaluated is whether AutoTIDBD is capable of performing representation learning.

## 8. Can AutoTIDB Perform Representation Learning?

We have demonstrated that AutoTIDBD is able to outperform scalar step size adaptation methods and ordinary TD on real-world prediction problems by tuning its step sizes, and that it is less sensitive to its parameters than ordinary TD. What remains of our four criteria is to determine if TIDBD is able to effectively perform representation learning by giving features relevant to the current prediction task large step sizes, and small step sizes to features which are irrelevant. We assess AutoTIDBD’s ability to perform representation learning, by analyzing the change in step size values for the prediction task introduced in Section 7 at  $\lambda = 0.95$ .

We created poor features by randomly choosing 25% of the features to be noisy. Noisy features were activated equiprobably. After completion of the experiment, the noisy features were compared with the number activations of each feature to ensure that noisy features included some which were highly active. There are certain time-steps for which all of the step sizes suddenly decrease. On these time-steps, the effective step size was greater than 1, leading to the normalization of the meta-weights to prevent over-shooting.

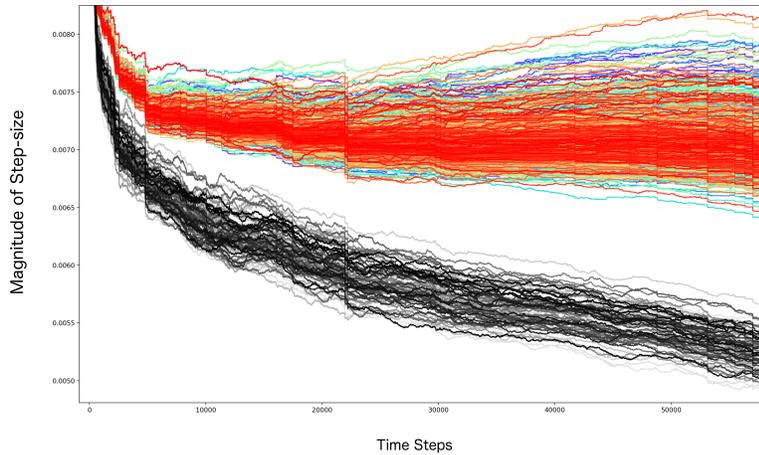


Figure 14: Average magnitude of step sizes over all trials. Noisy features are in greyscale, ordinary features are in colour.

Figure 14 depicts the magnitude of all step sizes averaged over all trials. step sizes corresponding to noisy features consistently decrease over time. We can see that the noisy features coloured in greyscale constantly shrink as experience determines them to be unreliable. This creates a separation between features which are noisy and those which are not. None of the noisy features have values within the range of ordinary features—all of the noisy features were correctly given smaller step sizes.

AutoTIDBD is able to perform representation learning by assigning appropriate step sizes, meeting our final criteria.

## 9. Limitations and Future Work

This derivation of TIDBD is limited to methods which use linear function approximation. Generalizations of IDBD for non-linear supervised learning methods exist (Schraudolph 1999), and could be generalized to TD learning in the future. In addition, this generalization is limited to on-policy predictions with replacing and accumulating traces. Further extension is required before AutoTDBD can be used with off-policy prediction methods and, and methods with different eligibility traces such as True Online TD (Seijen and Sutton, 2014).

In addition to producing more general methods, future research could pursue additional uses of learned step sizes. Step sizes learned with IDBD methods describe the relevance of given feature to the task at hand. If many of the features are large, then the features for the given prediction are well specified. This possibly be used to evaluate the potential of a prediction based on its given feature representation before it has been completely learned. Preliminary evaluation of predictions based on step sizes could be beneficial to prediction architectures such as horde (Sutton et al., 2011), where large collections of predictions proposed, learned, and maintained in real-time as a learner is interacting with their environment. In such situations limited computational resources must be used effectively; being able to better identify promising predictions in early learning could support prediction discovery in predictive knowledge systems.

Learned step sizes may also be an effective way to drive computational curiosity and intrinsic motivation. A challenge for learning systems is deciding how to explore their environment to support learning. Many intrinsic motivation systems rely on metrics which drive exploration based on error on a given task (Oudeyer and Kaplan 2009). One shortfall of these approaches is difficult for these methods to differentiate between situations where the error is high because not enough learning has occurred, and situations where the error is high because some signal or portion of the environment is not learnable. Learned step sizes describe how much learning, if used in combination with traditional error-based forms of intrinsic motivation, it may be better able to differentiate between what is novel and should be learned about , and what is unlearnable.

## 10. Conclusion

We presented an approach to generalizing Incremental Delta-Bar-Delta to temporal-difference learning, demonstrating that its effectiveness carries over from supervised learning to TD. We extended TIDBD to AutoTIDBD, using normalization methods from Autostep to improve the robustness of TIDBD. Adapting step sizes with AutoTIDBD is an improvement over ordinary TD methods with a tuned static step size, even on stationary problems. On non-stationary tasks, we showed that AutoTIDBD is able to find appropriate step sizes and differentiate between relevant and irrelevant features. Most importantly, over a number of real-world robotic prediction tasks we demonstrated that AutoTIDBD is less sensitive to choices of meta step sizes  $\theta$  and initial step sizes  $\alpha_0$  than ordinary TD is to settings of  $\alpha$ . AutoTIDBD out-performs TD for broad a broad range of meta step size settings which is relatively invariant over prediction problems. AutoTIDBD and TIDBD-based step size learning systems show promise of learning feature relevance and performing meta learning in an incrementally and online, lessening dependence on feature construction and parameter tuning.

## References

### References

- Barnard, E. (1993). Temporal-difference methods and Markov models. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(2):357–365.
- Colunga, E. and Smith, L. B. (2005). From the lexicon to expectations about kinds: A role for associative learning. *Psychological review*, 112(2):347.
- Dabney, W. and Barto, A. G. (2012). Adaptive Step-Size for Online Temporal Difference Learning. In *AAAI*.
- Dabney, W. C. (2014). *Adaptive Step-Sizes for Reinforcement Learning*. PhD thesis, University of Massachusetts Amherst.
- Dawson, M. R., Sherstan, C., Carey, J. P., Hebert, J. S., and Pilarski, P. M. (2014). Development of the Bento Arm: An improved robotic arm for myoelectric training and research. *Proceedings of MEC*, 14:60–64.

- Edwards, A. L., Hebert, J. S., and Pilarski, P. M. (2016). Machine learning and unlearning to autonomously switch between the functions of a myoelectric arm. In *Biomedical Robotics and Biomechatronics (BioRob), 2016 6th IEEE International Conference On*, pages 514–521. IEEE.
- Finn, C., Abbeel, P., and Levine, S. (2017). Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *arXiv:1703.03400 [cs]*.
- Kearney, A., Veeriah, V., Travnik, J., Sutton, R. S., and Pilarski, P. M. (2017). Every step you take: Vectorized adaptive step sizes for temporal difference learning. In *3rd Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, The University of Michigan, Ann Arbor, Michigan, USA.
- Mahmood, A. R. and Sutton, R. S. (2013). Representation Search through Generate and Test. In *AAAI Workshop: Learning Rich Representations from Low-Level Sensors*.
- Mahmood, A. R., Sutton, R. S., Degris, T., and Pilarski, P. M. (2012). Tuning-free step-size adaptation. In *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference On*, pages 2121–2124. IEEE.
- Oudeyer, P.-Y. and Kaplan, F. (2009). What is intrinsic motivation? A typology of computational approaches. *Frontiers in neurorobotics*, 1:6.
- Pilarski, P. M., Dawson, M. R., Degris, T., Carey, J. P., Chan, K. M., Hebert, J. S., and Sutton, R. S. (2013). Adaptive artificial limbs: A real-time approach to prediction and anticipation. *IEEE Robotics & Automation Magazine*, 20(1):53–64.
- Pilarski, P. M., Dawson, M. R., Degris, T., Carey, J. P., and Sutton, R. S. (2012). Dynamic switching and real-time machine learning for improved human control of assistive biomedical robots. In *Biomedical Robotics and Biomechatronics (BioRob), 2012 4th IEEE RAS & EMBS International Conference On*, pages 296–302. IEEE.
- Schraudolph, N. N. (1999). Local gain adaptation in stochastic gradient descent. In *9th International Conference on Artificial Neural Networks: ICANN '99*, pages 569–574.

- Schweighofer, N. and Arbib, M. A. (1998). A model of cerebellar metaplasticity. *Learning & Memory*, 4(5):421–428.
- Seijen, H. and Sutton, R. (2014). True online TD ( $\lambda$ ). In *International Conference on Machine Learning*, pages 692–700.
- Sutton, R. S. (1992). Adapting bias by gradient descent: An incremental version of delta-bar-delta. In *AAAI*, pages 171–176.
- Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT press Cambridge.
- Sutton, R. S., Modayil, J., Delp, M., Degris, T., Pilarski, P. M., White, A., and Precup, D. (2011). Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *AAMAS 2011*, pages 761–768. International Foundation for Autonomous Agents and Multiagent Systems.
- Travnik, J. B. and Pilarski, P. M. (2017). Representing high-dimensional data to intelligent prostheses and other wearable assistive robots: A first comparison of tile coding and selective Kanerva coding. *IEEE International Conference on Rehabilitation Robotics: [proceedings]*, 2017:1443–1450.
- Wilson, R. C. and Niv, Y. (2012). Inferring relevance in a changing world. *Frontiers in human neuroscience*, 5:189.
- Xu, Z., van Hasselt, H., and Silver, D. (2018). Meta-Gradient Reinforcement Learning. *arXiv:1805.09801 [cs, stat]*.
- Young, K., Wang, B., and Taylor, M. E. (2018). Metatrace: Online Step-size Tuning by Meta-gradient Descent for Reinforcement Learning Control. *arXiv:1805.04514 [cs, stat]*.